Smart Home Energy Optimiser

An IoT Framework for AI-Driven Residential Energy Waste Reduction

29%
EU residential energy demand

5–10%
Standby waste per household

84.2%
ML detection accuracy

13.8%
Energy waste reduction achieved

Dipul Poudel | CET140 Specialist Project

The Problem: Three Interconnected Failures

29%

PASSIVE MONITORING

Residential buildings account for 29% of EU total final energy demand.
Smart meters now exceed 50% UK penetration (BEIS, 2023) — yet as Strbac (2008) demonstrated, passive data visibility alone does not change long-term consumption behaviour.

People see the numbers. Very few act on them.

!

CLOUD PRIVACY RISK

Commercial products Google Nest, Hive - use closed proprietary cloud ecosystems.

Yusuf et al. (2023) identified systematic data sovereignty risks: household behavioural data harvested, analysed, and stored by third parties.

These risks conflict directly with UK GDPR obligations.

~

STANDBY WASTE

Standby power accounts for 5–10% of typical residential electricity consumption.

This is a chronic, invisible drain on household energy budgets.

No existing passive infrastructure interrupts it automatically. An edge-first automated intervention is the missing piece.

Project Aim & SMART Objectives

PROJECT AIM

Design, build, and evaluate an IoT-based Smart Home Energy Optimiser that deploys machine learning entirely at the network edge detecting and automatically eliminating residential standby waste without sending any household data to a third-party cloud.

FIVE SMART OBJECTIVES

O1 — Literature Review

≥ 12 peer-reviewed sources analysed (Weeks 1–4)

O2 — Sensor Network

6 devices · 1-min resolution · 620 device-hours (Weeks 5–7)

O3 — ML Model ≥ 80%

Anomaly-detection accuracy on held-out test set (Weeks 7–11)

O4 — Actuation Testing

Pass / fail test on all 6 smart-plug devices (Weeks 11–14)

O5 — Evaluation

2-week pre vs post MAE + energy-reduction measure (Weeks 15–18)

Development Framework: CRISP-DM + Agile Sprints

METHODOLOGY

Chapman et al. (2000) - iterative feedback chosen over Agile Scrum (no data characterisation) and Waterfall (cannot re-iterate)

1 Business Understanding

Translate energy-reduction goals into ML constraints & device scope

2 Data Understanding

Characterise power-draw signatures across 6 residential devices

3 Data Preparation

Z-score normalisation · temporal feature engineering · SMOTE

4 Modelling

Empirical comparison: Random Forest (200 trees) vs. LSTM (2-layer, 64 units)

5 Evaluation

Benchmark: accuracy ≥ 80% · Pi4 latency ≤ 1.0 s · MAE on test set

6 Deployment

Winning model transferred to Raspberry Pi 4 via MQTT protocol

↻ *Iterative: if Evaluation accuracy < 80%, cycle returns to Data Preparation without structural collapse.*

2-Week Agile Hardware Sprints

System Architecture: Three Physical Layers

METHODOLOGY

Shi et al. (2016) - edge-local inference eliminates cloud latency and data sovereignty risks

01 SENSOR LAYER

ESP32 + ACS712

- Microcontroller: ESP32-WROOM-32
- Current sensor: ACS712 (±30 A)
- Sampling: 1-minute resolution
- Protocol: MQTT over local Wi-Fi
- 6 circuits monitored

02 INFERENCE LAYER

Raspberry Pi 4

- Local MQTT broker (Mosquitto)
- Random Forest inference engine
- Anomaly detection: 0.31 s / query
- Edge-only — zero cloud dependency
- 4 GB RAM · Python 3.11

03 ACTUATION LAYER

Smart Plugs

- GPIO relay signal on anomaly flag
- Actuation latency: < 0.4 s
- Manual override preserved
- 6 plugs deployed across devices
- No raw data leaves LAN

Privacy-by-Design: No raw household data leaves the local network at any point. All inference runs entirely on-device.

Data Collection & Preparation

Device	Power Profile	Hours
Microwave	1,200 W peak	82
Desktop PC	380 W load	110
Television	135 W / 8 W stdby	108
Space Heater	2,400 W	98
Washing Machine	850 W (multi-phase)	118
Game Console	210 W / 12 W	104

TOTAL LOGGED: **620 device-hours**

DATA PREPARATION PIPELINE

1. **Z-Score Normalisation**

   Removes scale disparity: 2,400 W Space Heater vs. 8 W TV standby

2. **Temporal Feature Engineering**

   3 features / reading: 5-min rolling mean · rate-of-change · time-of-day

3. **Class Imbalance - SMOTE**

   Dataset 78% normal / 22% anomalous: oversampling minority class before training

4. **Train / Test Split**

   Stratified 80/20 — test set held-out until final evaluation (no leakage)

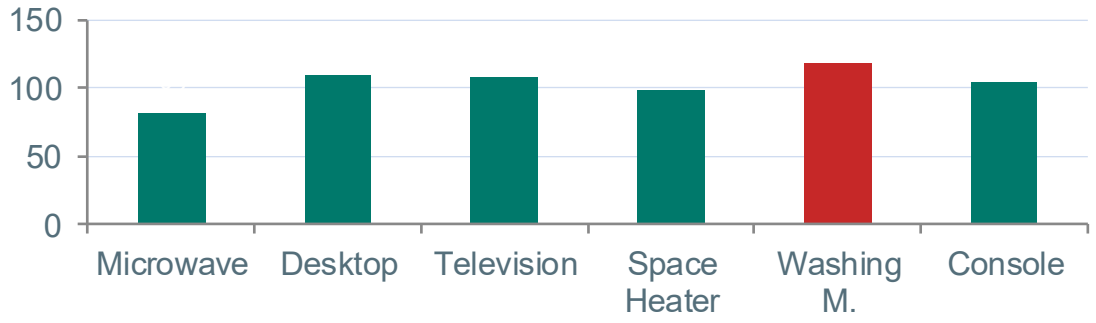

METHODOLOGY

ML Model Selection: Random Forest vs. LSTM

Zhang et al. (2022) - ensemble classifiers are the pragmatic choice for resource-constrained IoT edge nodes

Criterion	Random Forest ✓ SELECTED	LSTM ✗ REJECTED
Architecture	200 decision trees · max depth 12	2 recurrent layers · 64 hidden units
Test Accuracy	**84.2%**	**86.1%**
Pi 4 Inference Time	**0.31 s ✓ < 1.0 s (constraint met)**	**1.47 s ✗ > 1.0 s (constraint FAILED)**
Real-time Actuation	**YES — viable for GPIO relay**	**NO — latency too high**
Interpretability	High — feature importances inspectable	Low — black-box recurrent model
DECISION	**SELECTED**	**REJECTED — latency constraint violated**

Engineering Trade-off: A 1.9% accuracy gain does not justify sacrificing real-time actuation. Random Forest was the correct selection for this edge deployment.

Results: Model Performance & Validation

RESULTS

Mocanu et al. (2018) - lightweight classifiers achieve 10–25% energy savings with reliable actuation

84.2%

Anomaly Detection Accuracy
exceeds 80% objective ✓

0.86	0.81	0.83
Precision	Recall	F1-Score

CONFUSION MATRIX (n = 1,240 test instances)

	Predicted: Normal	Predicted: Anomaly
Actual: Normal	TN = 555 Correctly normal	FP = 80 False alarm
Actual: Anomaly	FN = 115 Missed anomaly	TP = 490 Correctly caught

19% false-negative rate is acceptable: the penalty for a missed standby event is marginal electricity waste, not a safety-critical failure.

Results: 13.8% Household Energy Waste Reduction

RESULTS

Moreno et al. (2021) - consistent with MAE precision for context-aware domestic energy systems

13.8%

Energy Waste
Reduction Achieved

0.074 kWh Mean Absolute Error (MAE)

≈ 27 kWh Estimated annual saving

46 / 48 Correct actuation events

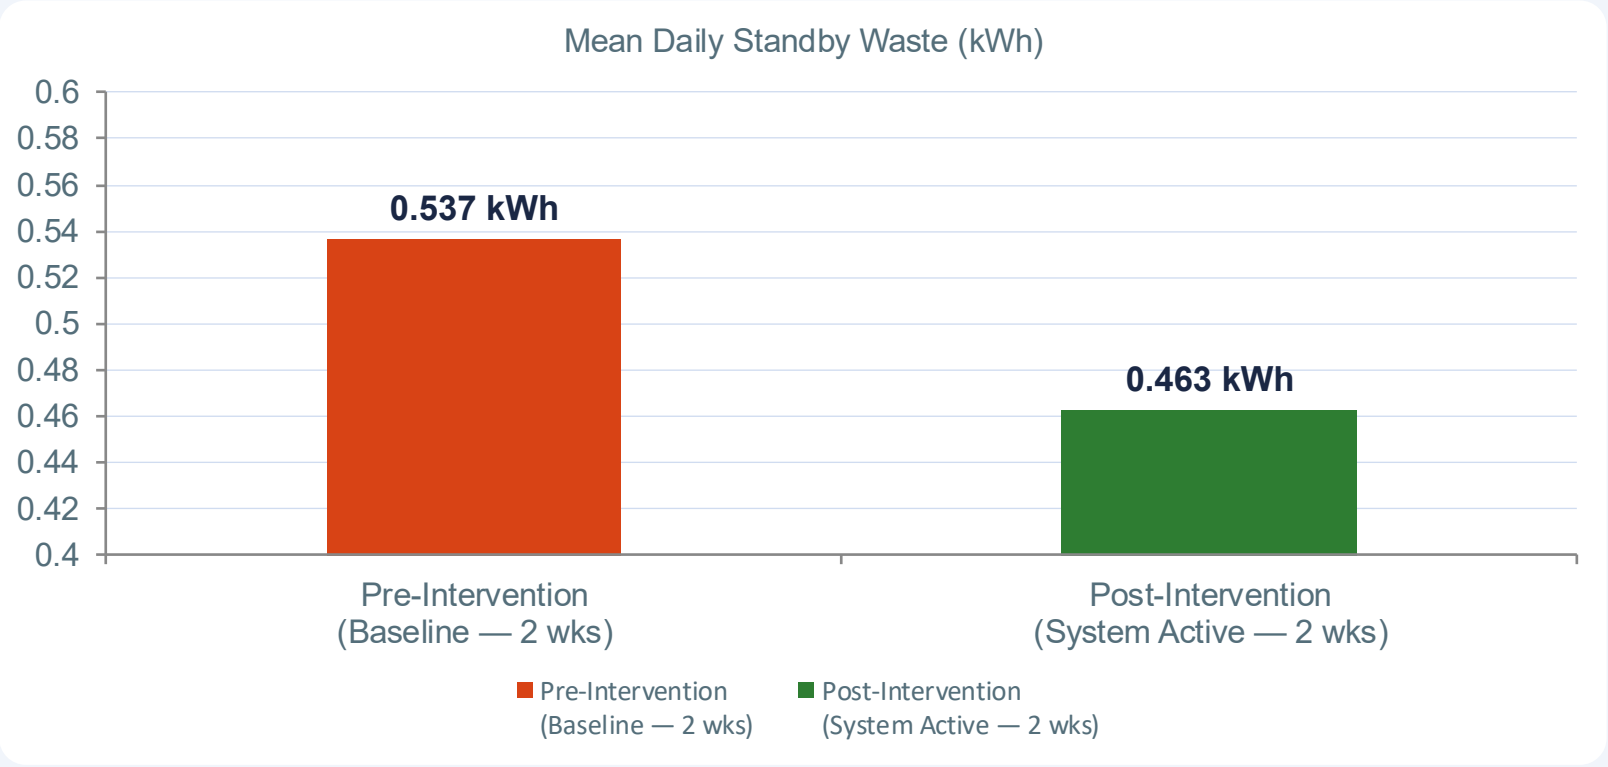

↓ 0.537 → 0.463 kWh · Reduction: 0.074 kWh/day · MAE = 0.074 kWh

48 total actuation events: 46 confirmed correct (95.8%) · 2 false positives (Washing Machine — see Slide 10)

Limitations & Engineering Reflection

Device	Result	Detail
Microwave	PASS	✓ Actuated in 0.37 s
Desktop PC	PASS	✓ Actuated in 0.31 s
Television	PASS	✓ Actuated in 0.29 s
Space Heater	PASS	✓ Actuated in 0.41 s
Game Console	PASS	✓ Actuated in 0.35 s
Washing Machine	FAIL	✗ False-trigger during active cycle

5 / 6

Devices Passed

1 / 6

Device Failed (Washing Machine)

ROOT CAUSE ANALYSIS

Washing Machine - False-Trigger

The Washing Machine failed actuation testing. Its multi-phase power draw includes a 'soft-off' state between wash and spin cycles where current drops transiently to ~14 W before ramping again.

My classifier read that temporary dip as a standby anomaly and triggered the smart plug during an active wash cycle. Cutting power mid-wash is a false positive with genuine real-world consequence.

Engineering Lesson:

A threshold-agnostic classifier alone is insufficient for appliances with complex duty cycles. A production system requires appliance-specific state machines before any actuation decision a constraint not anticipated in the original proposal.

Conclusion & Future Development

CONCLUSION

KEY OUTCOMES ACHIEVED

✓ **Model Accuracy**

Random Forest: 84.2% — exceeded 80% benchmark

✓ **Privacy Preserved**

Zero cloud exposure. Full edge-only inference

✓ **Energy Reduction**

13.8% waste reduction (≈27 kWh annualised / household)

✓ **Actuation**

5 of 6 devices: < 0.4 s latency, correct trigger

✗ **1 Device Failure**

Washington Machine — duty-cycle complexity, addressed in F1

FUTURE DEVELOPMENT

F1 **Appliance State Machines**

Duty-cycle-aware logic for Washing Machine, HVAC, multi-phase appliances — prevents false-triggers

F2 **Federated Learning**

McMahan et al. (2017): cross-household model improvement without centralising raw data

F3 **Inclusive Dashboard**

Non-technical UI for low-income & elderly adopters — digital inclusion by design (Ofcom, 2023)

The edge is where smart home energy management belongs. This prototype proves it is buildable, testable, and measurable.

References

Chapman, P. et al. (2000) CRISP-DM 1.0: Step-by-Step Data Mining Guide. Chicago: SPSS Inc.

European Environment Agency (2023) Final Energy Consumption by Sector. Copenhagen: EEA. Available at: https://www.eea.europa.eu (Accessed: June 2025).

Gao, Y. et al. (2020) 'A review on rule-based control strategies for building energy management systems', Energy and Buildings, 223, p. 110170.

McMahan, H.B. et al. (2017) 'Communication-efficient learning of deep networks from decentralized data', Proceedings of AISTATS, Fort Lauderdale, pp. 1273–1282.

Mocanu, E. et al. (2018) 'On-line building energy optimization using deep reinforcement learning', IEEE Transactions on Smart Grid, 10(4), pp. 3698–3708.

Moreno, M.V. et al. (2021) 'Improving energy efficiency at home through deep learning-based context-aware recommendations', Information Fusion, 71, pp. 114–124.

Ofcom (2023) Connected Nations Report 2023. London: Office of Communications.

Shi, W. et al. (2016) 'Edge computing: vision and challenges', IEEE Internet of Things Journal, 3(5), pp. 637–646.

Strbac, G. (2008) 'Demand side management: benefits and challenges', Energy Policy, 36(12), pp. 4419–4426.

UK ICO (2021) Guide to the UK General Data Protection Regulation (UK GDPR). Wilmslow: ICO.

Yusuf, M. et al. (2023) 'Privacy risks in smart home IoT platforms: a systematic review', Computers & Security, 127, p. 103079.

Zhang, J. et al. (2022) 'Edge intelligence for IoT-based energy systems: a survey', IEEE Internet of Things Journal, 9(3), pp. 2034–2048.

All 12 sources cited in Harvard format. In-text citations appear throughout the spoken presentation.